

UCRL-JC-124245

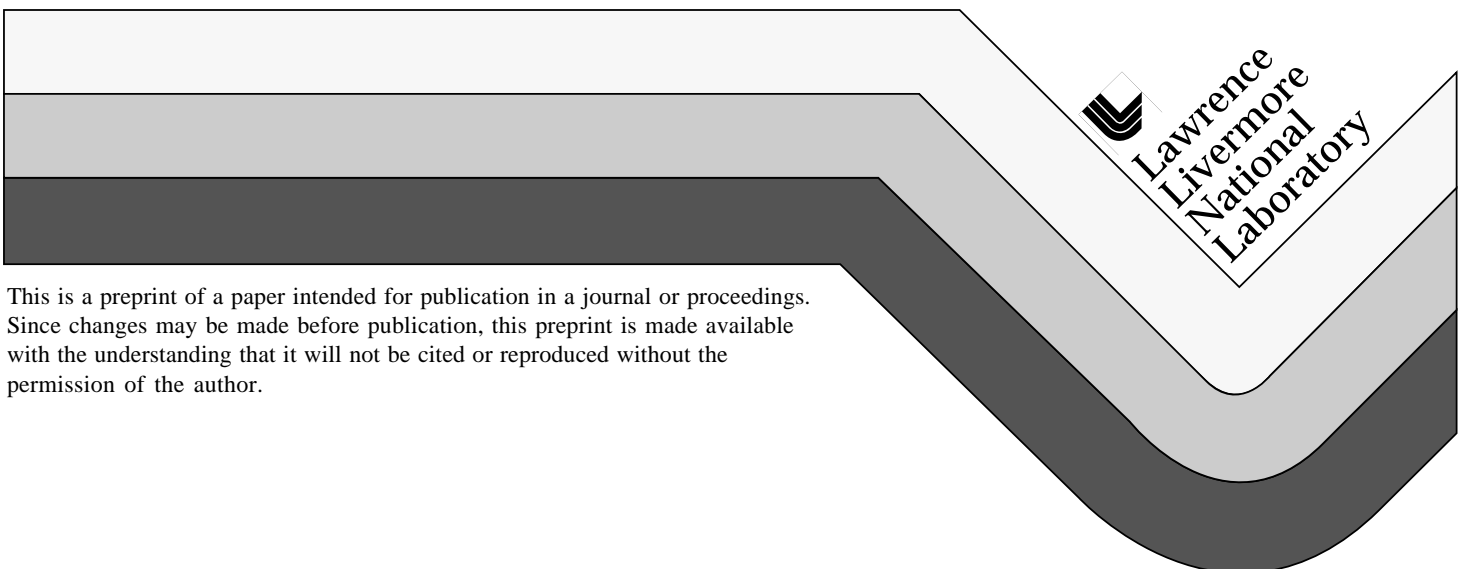
PREPRINT

Large-Timestep Techniques for Particle-In-Cell Simulation of Systems with Applied Fields that Vary Rapidly in Space

A. Friedman
D. P. Grote

This paper was prepared for submittal to the
American Institute of Physics Conference Proceedings
Computational Accelerator Physics Conference
Williamsburg, VA
September 23-27, 1996

October 3, 1996



This is a preprint of a paper intended for publication in a journal or proceedings.
Since changes may be made before publication, this preprint is made available
with the understanding that it will not be cited or reproduced without the
permission of the author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Large-Timestep Techniques for Particle-In-Cell Simulation of Systems with Applied Fields that Vary Rapidly in Space*

Alex Friedman and David P. Grote

*Lawrence Livermore National Laboratory
L-440, P.O. Box 5508, Livermore CA 94550*

Abstract. Under conditions which arise commonly in space-charge-dominated beam applications, the applied focusing, bending, and accelerating fields vary rapidly with axial position, while the self-fields (which are, on average, comparable in strength to the applied fields) vary smoothly. In such cases it is desirable to employ timesteps which advance the particles over distances greater than the characteristic scales over which the applied fields vary. Several related concepts are potentially applicable: sub-cycling of the particle advance relative to the field solution, a higher-order time-advance algorithm, force-averaging by integration along approximate orbits, and orbit-averaging. We report on our investigations into the utility of such techniques for systems typical of those encountered in accelerator studies for heavy-ion beam-driven inertial fusion.

I. INTRODUCTION

We are exploring techniques for enhanced efficiency in PIC simulations of beams and plasmas. A simple such technique, used in the WARP code[1,2] from its inception, is the “residence correction,” whereby the impulse imparted by an (idealized) sharp-edged element is corrected to reflect the fraction of the velocity-advance step during which the particle resides within the element. It is desirable to extend this concept to cases where the effects of extended fringe fields and other smooth but rapid variations must be accurately captured. Techniques which may be applicable include: a) a subcycled leapfrog advance—taking N substeps between each major step on which the self-consistent field is computed and applied; b) a family of high-order symplectic advances—for smooth forces, these converge quickly as the step size is reduced, but it will be nontrivial to take advantage of that property in a practical PIC code; c) a newly-invented family of force-averaged velocity advances—effectively, these amount to subcycling on the velocity advance only; using integration along an approximate orbit to do the averaging; however, “special” weightings may afford higher order, and again the challenge is to make the methods practical; and d) orbit-averaging—this is basically a noise-reduction strategy which may allow use of a smaller number of particles, and probably must be non-symplectic; it is discussed briefly here but has not been studied in detail. In Section II immediately below these methods are described. A model problem was devised to evaluate the methods; the problem and some results are described in Section III. We are beginning to test the most promising of these methods via full 3-d PIC simulations using the WARP3d code; some initial results, and a concluding discussion, are presented in Section IV.

II. METHODS CONSIDERED

A. Subcycled Leapfrog

In this algorithm, the external force is applied every substep, while a space charge “kick” is applied every N substeps, after the source term is accumulated and the field equation solved. It is necessary to apply the self-force only at those times at which the source is calculated, to avoid spurious self-forces, unless special steps are taken as in orbit-averaging. In the absence of space charge the subcycled advance is identical to ordinary leapfrog with a step size $1/N$ as large.

B. High-Order Symplectic Integrators

These methods preserve Hamiltonian structure and avoid spurious damping or excitation, but are not energy conserving.[3] Nonrelativistically, defining the position “ x ”, velocity “ v ”, and acceleration “ a ”, the general explicit algorithm is:

$$x' = x + c_i \Delta t v; \quad v' = v + d_i \Delta t a; \quad \text{repeated for } i = 1, \dots, k. \quad (1)$$

The c_i and d_i are chosen so that the final composed mapping satisfies the Taylor series expansion of the solution up to order Δt^n , for an n -th order scheme. The Candy-Rozmus (C-R) scheme has the advantage of offering 4th-order accuracy with 3 force evaluations per step. We suspect that rotations of v in an external magnetic field applied at each substep won't break the invariants, since they are 1:1 maps, but have yet to prove this. Coefficients for some schemes are:

$$\text{Leapfrog:} \quad c_1 = 0, \quad c_2 = 1, \quad d_1 = d_2 = 1/2 \quad (2)$$

$$\text{Ruth 3rd-order:} \quad c_1 = 7/24, \quad c_2 = 3/4, \quad c_3 = -1/24 \quad (3) \\ d_1 = 2/3, \quad d_2 = -2/3, \quad d_3 = 1$$

$$\text{Candy \& Rozmus:} \quad c_1 = c_4 = 1/[2(2-2^{1/3})], \quad c_2 = c_3 = [1-2^{1/3}]/[2(2-2^{1/3})] \quad (4) \\ d_1 = d_3 = 1/(2-2^{1/3}), \quad d_2 = -2^{1/3}/(2-2^{1/3}), \quad d_4 = 0$$

C. Force-Averaged Velocity Advance

The original concept was to integrate the external force over the velocity-advance step along an approximate orbit, for an improved impulse. This generalizes the “residence correction” used in WARP3d on entry to or exit from a sharp-edged element. The averaged applied force is: $\langle F_{\text{ext}} \rangle = \sum W_j F_{\text{ext}}(x_j)$, where the x_j are computed at temporal offsets δ_j relative to the middle of the velocity-advance step, and the W_j are “weights.” Letting “sc” denote space charge, $*$ a temporary quantity, “h” the half-level, and acceleration $a = F/m$, a timestep is:

$$v_h := v + 1/2 \Delta t \langle a \rangle \\ x := x + \Delta t v_h \\ a^* := a_{\text{ext}}(x) + a_{\text{sc}}(x)$$

$$\begin{aligned}
v^* &:= v_h + 1/2 \Delta t a^* \\
\langle a_{\text{ext}} \rangle &:= 0 \\
\text{for } 1 \leq j \leq k : \quad &x_j := x + v^* \delta_j + 1/2 a^* \delta_j^2 ; \\
&\langle a_{\text{ext}} \rangle := \langle a_{\text{ext}} \rangle + W_j a_{\text{ext}}(x_j) \\
\langle a \rangle &:= \langle a_{\text{ext}} \rangle + a_{\text{sc}}(x) \\
v &:= v_h + 1/2 \Delta t \langle a \rangle
\end{aligned} \tag{5}$$

During the evaluation of force-averaged methods using the model problem, it was discovered that certain simple quadratures work “better” than a detailed integration using many points (very large k , all W ’s equal). One “good” choice, found via experimentation for $k = 3$, is: $W = \{1/3, 1/3, 1/3\}$, $\delta = \Delta t \{-1/2, 0, 1/2\}$. The simplest “good” scheme, found via the analysis outlined below for $k = 2$, is: $W = \{1/2, 1/2\}$, $\delta = \Delta t\{-1/\sqrt{6}, 1/\sqrt{6}\}$. However, it is difficult to start the calculation so as to preserve accuracy, and we have not yet learned how to do so reliably. The difficulty arises when the trajectory calculation is initiated at an axial position where the applied force has a nonzero axial gradient.

To analyze the force-averaged velocity advance, we assumed harmonic motion at frequency ω in a well of natural frequency ω_0 , *i.e.*, $x^{n+1} = e^{-i\omega\Delta} x^n$ and $a^n = -\omega_0^2 x^n$ (superscripts denote time levels). We also insist that $\sum W_j = 1$, $\sum \delta_j = 0$, and $\sum W_j \delta_j = 0$ (W ’s symmetric about the middle of the set, δ ’s antisymmetric). Computing ω in terms of ω_0 , Δt , W_j , and δ_j , one obtains the condition (satisfied by the “good” cases mentioned above, and others) for fourth-order agreement between ω and ω_0 , *i.e.*, $\omega / \omega_0 = 1 + O(\omega\Delta)^4$:

$$\sum W_j \delta_j^2 = (\Delta t)^2 / 6 . \tag{6}$$

D. Orbit-averaging

In orbit averaging, the field equation is solved every N substeps using a net charge that comes from a sum over depositions done each substep; the field is applied every substep using interpolation in time, in a predictor-corrector loop. In general, the substep will be set by the need to resolve external field gradients. In comparison with leapfrog, there can be a gain in efficiency due to less-frequent field-solving. In comparison with subcycling, there can be reduced self-field noise, at the expense of extra deposition time and the need for predictor-corrector iteration. A disadvantage is that such methods are unlikely to be symplectic, but this is probably not a problem for our applications, and in any event results can be validated by varying the step size. Recent research in France exploring this technique has been described as encouraging [4]

III. MODEL PROBLEM

This problem models transport through a lattice period (a pair of quadrupole lenses of alternating polarity), tracking a particle on the principal axis $y = 0$.

Linear fringe fields ($\propto x$) are included, but higher multipoles and pseudo-multipoles are not included. “Exact” solutions (needed for the computation of errors) were obtained by use of very small steps in a leapfrog advance. The error for each run was defined as the absolute value of the difference between the “exact” and computed x values at the end of the problem. Three elements were varied, making a series of 12 tests. In each test, 5 algorithms were run at 7 step sizes ranging from 16 to 1024 steps per lattice period (256 is a typical value used in WARP3d leapfrog runs). The elements varied were:

(1) Three models of the force profile (shown in Fig. 1):

“sine⁵”: $F(x,z) = -x F_0 \sin^5 [(z + z_{\text{off}}) \pi/L]$

“model”: $F(x,z) = \pm 1/2 x F_0 \{ \cos[(z + z_{\text{off}}) \pi/\chi] + 1 \}$, or 0, pieced together

“tabulated”: Linear interpolation into a table of values derived analytically




FIGURE 1. Force profiles and trajectories, $z_{\text{off}}/\text{hlp}=0.5$

Here, L is the half-lattice period length, and χ the physical magnet length. Of these, “sine⁵” is perfectly smooth, while “model” has a discontinuous second derivative, and “tabulated” has a discontinuous first derivative due to abrupt termination of the list of 300 tabulated values describing each “peak,” which were offset and rescaled slightly so that the last tabulated force value was zero.

(2) Two options for modeling space charge: $\text{spacechg}=0$, no space charge force; and $\text{spacechg}=1$, a linear kick applied at end of each full cycle of the advance.

(3) Two offsets of the initial position relative to the center of a “drift space”: $z_{\text{off}}/\text{hlp}=0.5$, with the test particle launched from the center of a focusing quad; and $z_{\text{off}}/\text{hlp}=0.4$, with the particle launched from the “side of the hill.”

Some results are shown in Fig. 2. In (a), the “canonical” smooth sine⁵ case with the particle launched from the center of a quad and no space charge, the subcycled leapfrog is identical to ordinary leapfrog with a 4x smaller step, while

Ruth's scheme converges cubically, and both the Candy-Rozmus and the force-averaged scheme converge quartically. When (b) the tabulated force is used, all of the improved schemes do better than leapfrog, but none clearly converges more rapidly than quadratic; we conclude that a tabulated force with the granularity typically used is insufficiently smooth with linear interpolation, and expect that a cubic spline interpolation is needed for full benefit from the high-order schemes. When (c) the sine⁵ force is used with the particle launched from the side of the hill, the leapfrog error is considerably worsened, and the utter failure of the optimized force-averaged scheme is evident. The high-order schemes do not degrade the way leapfrog does. In this case the high-order advance performed so well that the "exact" answer was not quite "exact enough"; the quartic convergence of C-R actually persists down to the 1024-steps case. When (d) the tabulated force is used with a side-of-the-hill start, none of the high-order schemes does significantly better than subcycled leapfrog.

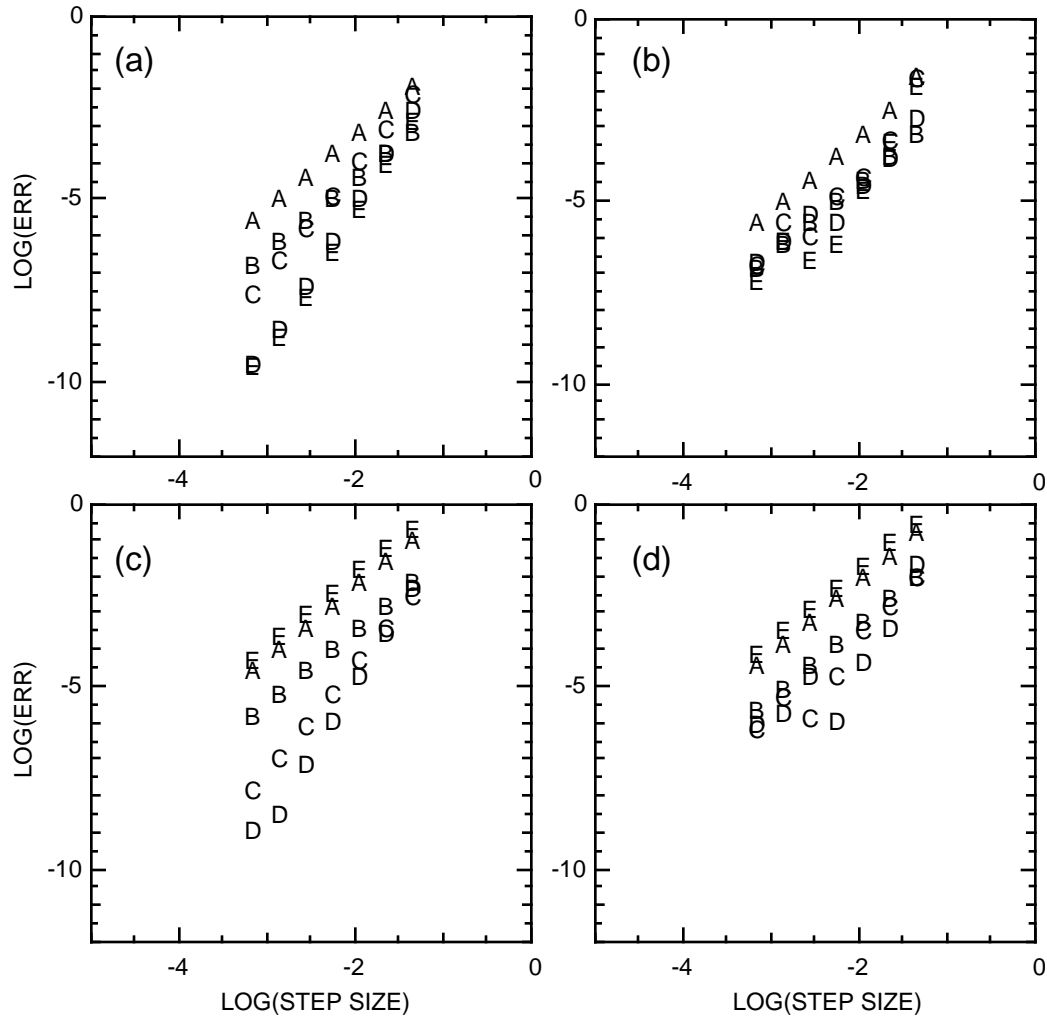


FIGURE 2. Error vs. step size: (a) $z_{\text{off}}/h_{\text{lp}}=0.5$, sine^5 force; (b) $z_{\text{off}}/h_{\text{lp}}=0.5$, tabulated force; (c) $z_{\text{off}}/h_{\text{lp}}=0.4$, sine^5 force; (d) $z_{\text{off}}/h_{\text{lp}}=0.4$, tabulated force. Points marked A were computed with leapfrog, B with 4:1 subcycled leapfrog, C with the Ruth scheme, D with the Candy-Rozmus scheme, and E with the optimized two-point force-averaged scheme.

When the sine^5 force is used with (non-subcycled) space charge kicks turned on (not shown), subcycled leapfrog is not as good as leapfrog with a 4x smaller step. The high-order schemes do somewhat better than subcycled leapfrog, but are ultimately just second order because the space-charge error eventually dominates. When the “model” force is used with $z_{\text{off}}/h_{\text{lp}}=0.4$, the Ruth scheme does better than the Candy-Rozmus one, and is anomalously fourth-order. Leapfrog does poorly on that problem as well. A number of other interesting features are evident in the runs, but space does not permit their explication here.

IV. SIMULATIONS AND DISCUSSION

Experiments at LLNL studying the bending of space-charge-dominated beams are simulated using WARP3d, a 3-d particle-in-cell code developed for heavy-ion fusion accelerator studies. The beamline includes seven electrostatic quadrupoles, eight permanent magnet quads, and a lattice of five electric dipoles interleaved with five permanent magnet quadrupoles. This serves as a realistic test of new methods. To date, 5:1 subcycling has been tested and seems to work well, with little distinguishable difference in the output. The speed-up in the case tested was relatively minor because a fast field-solver and a large particle number were used.

The utility of high-order advances is application-dependent. Such methods are often used in the absence of strong space-charge, and work well for the applied fields in typical applications. However, in a PIC code with multilinear interpolation (such as WARP3d), the requisite force smoothness may be absent—the pairwise interparticle force is continuous as a function of interparticle separation, but its derivative is “almost always” discontinuous, [5] and we don’t yet know if multilinear interpolation is “smooth enough” in practice. If it is not, a smooth, higher-order spatial interpolant could be used. The cost of interpolation would be increased, but this is likely to be usually unimportant. Since a finer grid would be undesirable in 3-D, one should avoid enlarging the effective particle size. “Sharpening” operators can be used to offset any spreading. [6] For the 4th-order advance, a three times larger step size would be needed for net gain if the self-consistent field were computed at each substep. However, in many problems that field might be computed and applied only once per step, as in subcycling.

We conclude that there are clear advantages to using one or more of these methods to model space-charge-dominated beams. Subcycling is simple and seems to work well, while the payoff from a sophisticated scheme may ultimately be larger. The “startup” question needs to be addressed. We plan to explore these options systematically, while taking advantage of what we have already learned.

REFERENCES

*Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-48.

1. A. Friedman, D. P. Grote, and I. Haber, *Phys. Fluids B* 4, 2203 (1992).
2. D. P. Grote, A. Friedman, I. Haber, and S. S. Yu, in *Proc. Int. Sympos. on Heavy Ion Inertial Fusion*, Princeton, Sept. 6-9, 1995; to be publ. in *Journal of Fusion Eng. Design*, 1996.
3. H. Yoshida, *Celestial Mech. and Dynamical Astron.* 56, 27-43 (1993).
4. A. Piquemal, private communication, June 1996.
5. C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation*, New York, McGraw-Hill, 1981, 70-72.
6. R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles*, Bristol, Adam Hilger, 1988, 137-140.